National eResearch Collaboration Tools and Resources

www.nectar.org.au

NeCTAR is an Australian Government project conducted as part of the Super Science initiative and financed by the Education Investment Fund. The University of Melbourne has been appointed the lead agent by the Commonwealth of Australia, Department of Innovation, Industry, Science and Research.

Objectives: to enhance research collaboration through the development of eResearch infrastructure.

Overall Timeline - Infrastructure Extension

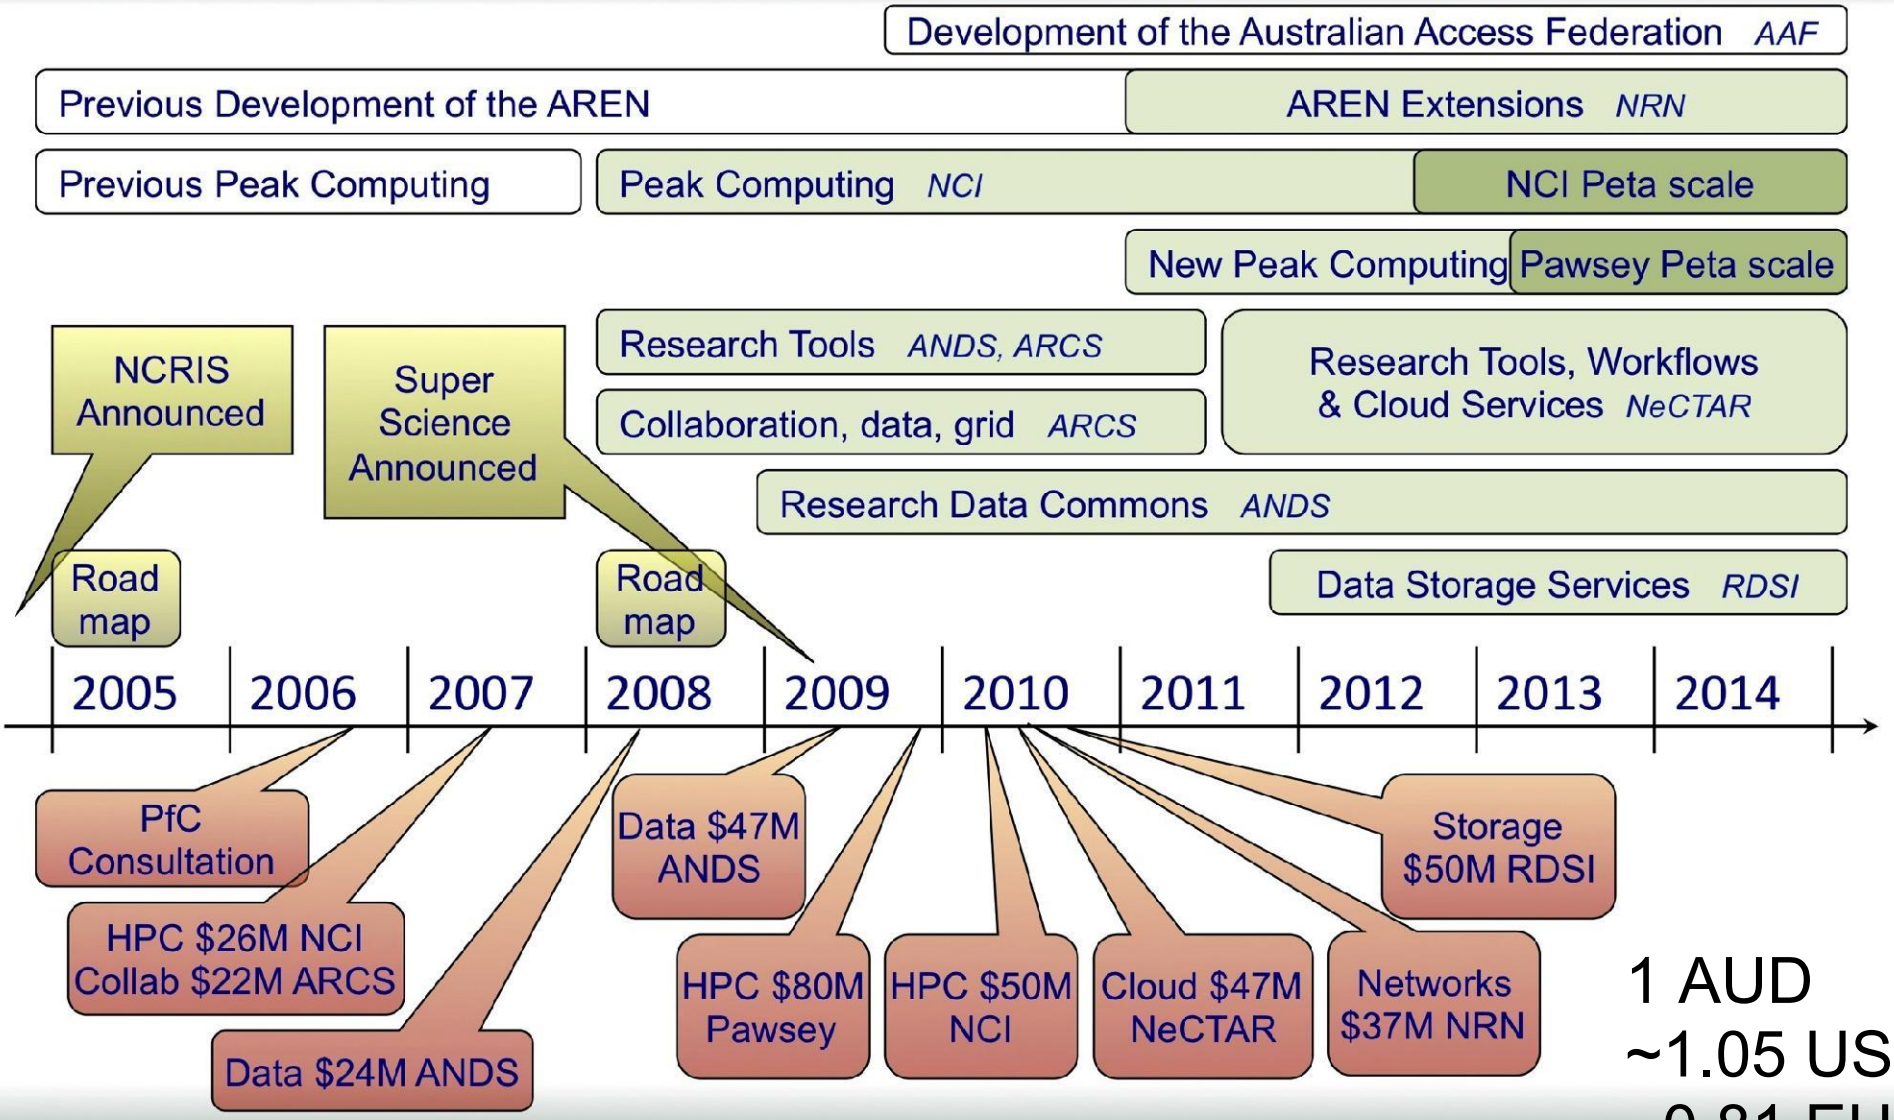

1 AUD
~1.05 US
~0.81 EU

NeCTAR has four program areas

Research software

Virtual laboratories

eResearch Tools

Computational platforms

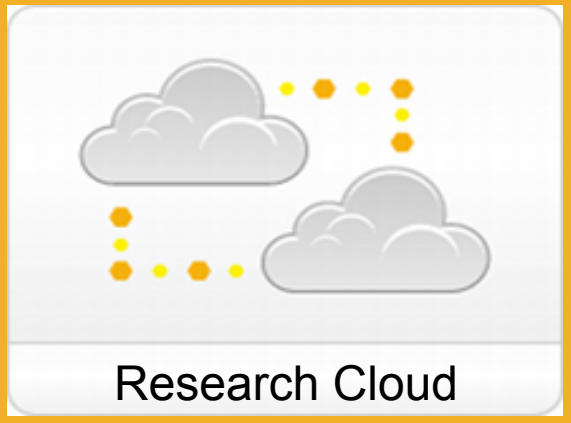

Research Cloud

National Server Program

nectar

Who we are - OpenStack Context

Cells, Database migrations, S3/EC2 APIs, code typos, HTTP/S headers, CA certificates, rabbit HA, cinder driver hacks,...

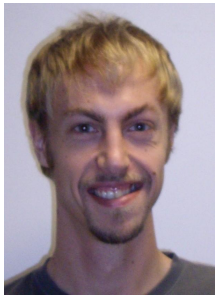

Horizon Core.
test fixes, sec groups, volume quotas, docs, usability, roles, cosmetics, bugs, pagination, permissions, ...

Documentation Core.

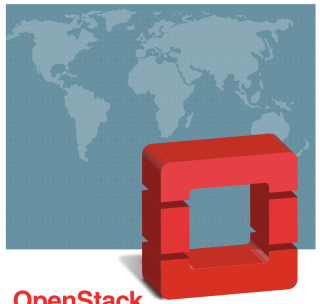

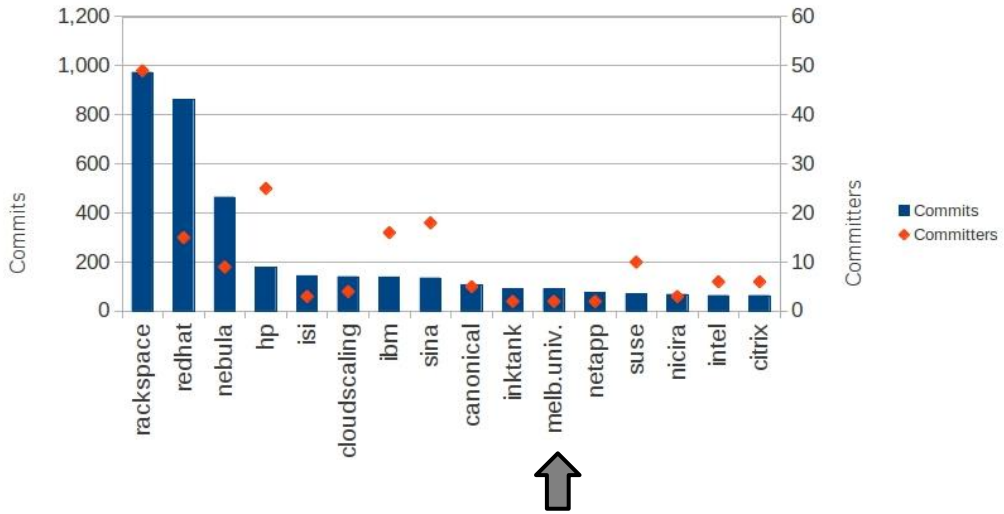

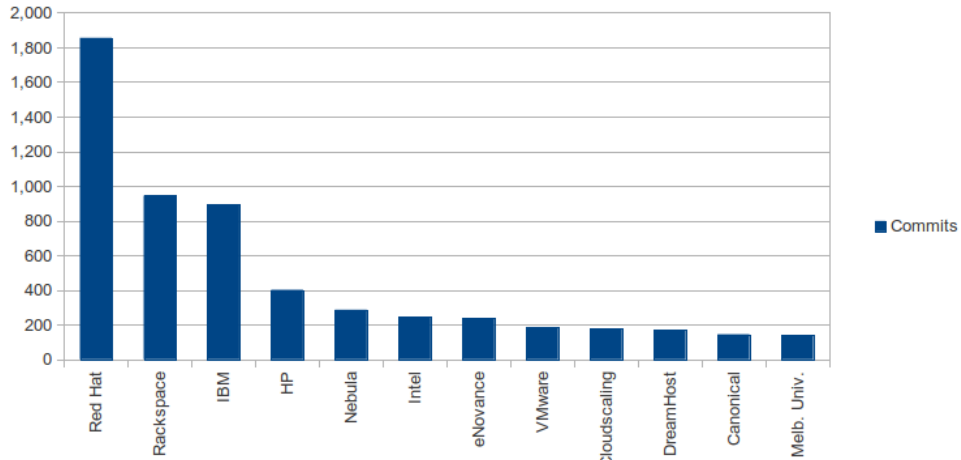

The NeCTAR Research Cloud

A platform for hosting, deploying and sharing research software infrastructure

Build to a research spec

- researchers work 24/7, globally, collaboratively across boundaries

An OpenStack cloud

- Any Researcher, Any Discipline, Anywhere

A single national cloud

- 8 sites, ~4k cores/site -> 30k cores

nectar

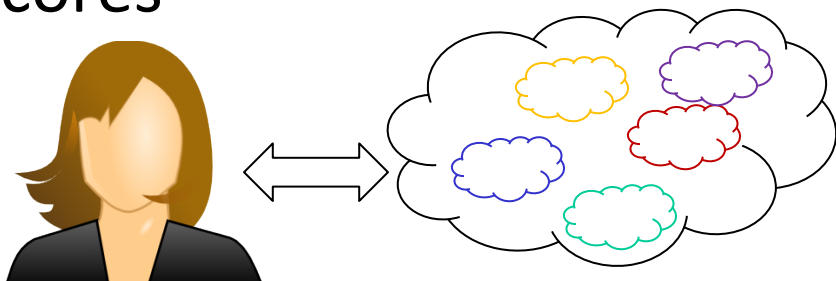

This cloud... why build it ourselves?

1. Proximity – the honeypot – infrastructure attracts community

2. Local infrastructure is more responsive to research needs

3. Service offering and usage modes suitable for research

4. Locality to instruments, research networks and other infrastructure.

5. Data sovereignty

nectar

nectar

New Australian research cloud, ready to use now

January 2012

April 2013

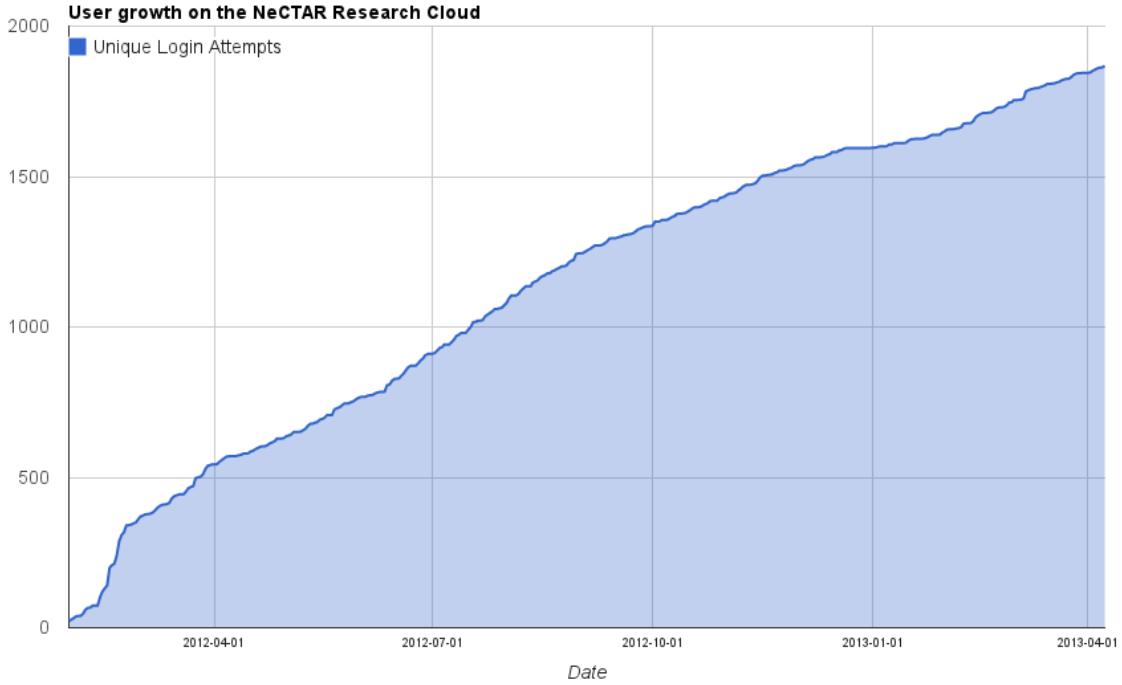

Hardware – in rack, right now

The University of Melbourne
336 cores – 48 core Dell R815s
3840 cores – 160 x 24 Core, 128GB, 10Gbit/s Xenon Quad2U
195TB – HP DL180G6 w/ DL2000 @ 24TB/node
146TB – Dell R715 w/ MD1200 @ 24TB/node
100TB – Hitachi HNAS/BlueARC
10Gbit/s – CISCO Nexus (2232, 5596, …)

Queensland Cyber Infrastructure Foundation
512 cores – 64 core, 256GB, SGI H2106-G7
180TB – SGI 3112 @ 36TB/node
10Gbit/s – CISCO Nexus (2232, 5548, …)

Monash University
2304 cores – 96 x 48 core, 192GB, 10Gbits, Dell C6145
192TB – Dell R720xd @ 24TB/node
10Gbit/s – CISCO Nexus (2248, 5596, …)

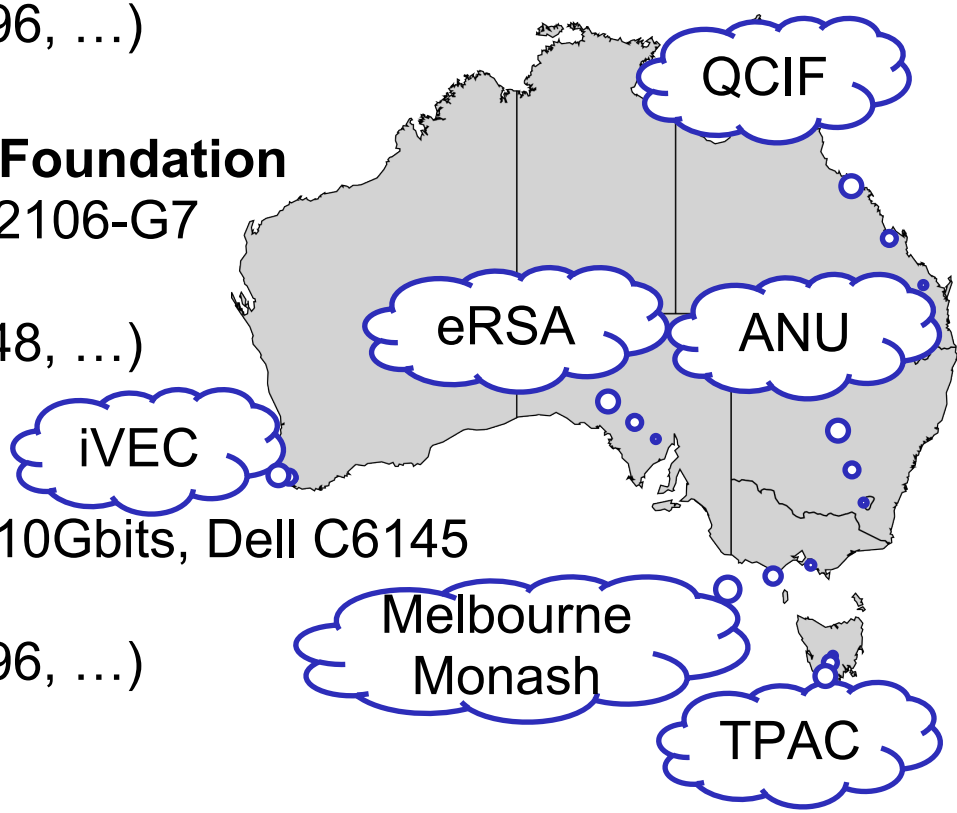

nectar

What we found – use cases

One core webserver

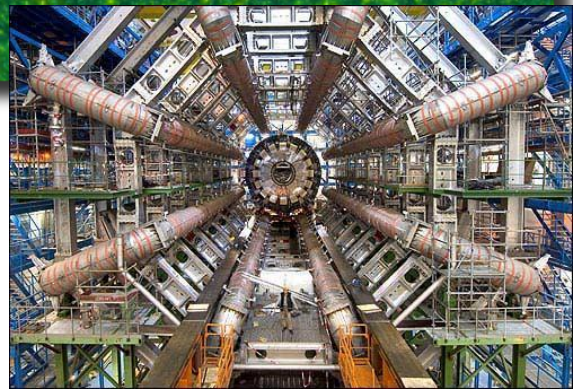

nectar

High Throughput

OzTrack

We're live - what problems have they had?

- Forgot to add a keypair, forgot to add an SSH group
- Too many security groups - can't see launch button
- Difficult to distinguish "good" images from "other"
- Don't understand storage types
- Client tools in official repositories too old
- S3 API not compatible
- Snapshots very slow to create
- Storage performance not good enough
- Users trying to make their own images, and failing

We're live - security incidents

Open DNS resolvers

1 spam source - SMTP running on port 80

User thought our security guidelines (no mailservers!) were too onerous and started arguing

1 spam source

A student was running the machine

1 Compromised machine clicking on web advertisements

nectar

We're live - What works? What doesn't?

- API scales out nicely
- Underlying Storage - not enough disks
- Object Storage works well, but has poor uptake
- Staffing skills
- Lack of operational tools takes up support time

nectar

Moving to Platforms

- NeCTAR has funded VLs and RTs, but so many more
- Researchers should not have to be sysadmins

Leverage capability from sector's organisations to support deployment

- Developer days around the country
- Focus on tapping into existing applications
- Within-institute, within-domain

Discoverability

- Not the "App Store" for the Cloud, but making it easy to share applications with colleagues
- Metadata about apps

nectar

Moving to Platforms

Make the cloud easy to use through use of

- Recipes
- Toolkits
- Scripts

Aim for migration from single-server app to dynamically scalable cloud app

The conversations we're having in 2012 (cloud) will be very different in 2013 (research) – Steven Manos

nectar

Why are we different - multiple companies?

- Each site runs:
  - Their own compute nodes inside 1 or more cells
  - Swift cluster
  - Glance-api node(s)

- Central Services
  - Keystone
  - glance-registry
  - Dashboard

- Resource allocation, researchers don't pay in $$

nectar

Why are we different? Openstack mods

- Try and run as close to stable releases as possible
- Shibboleth login to dashboard
- Geo Distributed glance-api's with central glance-registry. V2??

- Cells (more later)

nectar

Why are we different - puppet

- Develop all our own puppet modules
  - puppetlabs modules weren't around when we started
- All sites use a central puppet server and split up using puppet environments
- Use gerrit and Jenkins for QA
- https://github.com/NeCTAR-RC

Cells - explanation

Like Availability Zones in Amazon-speak?

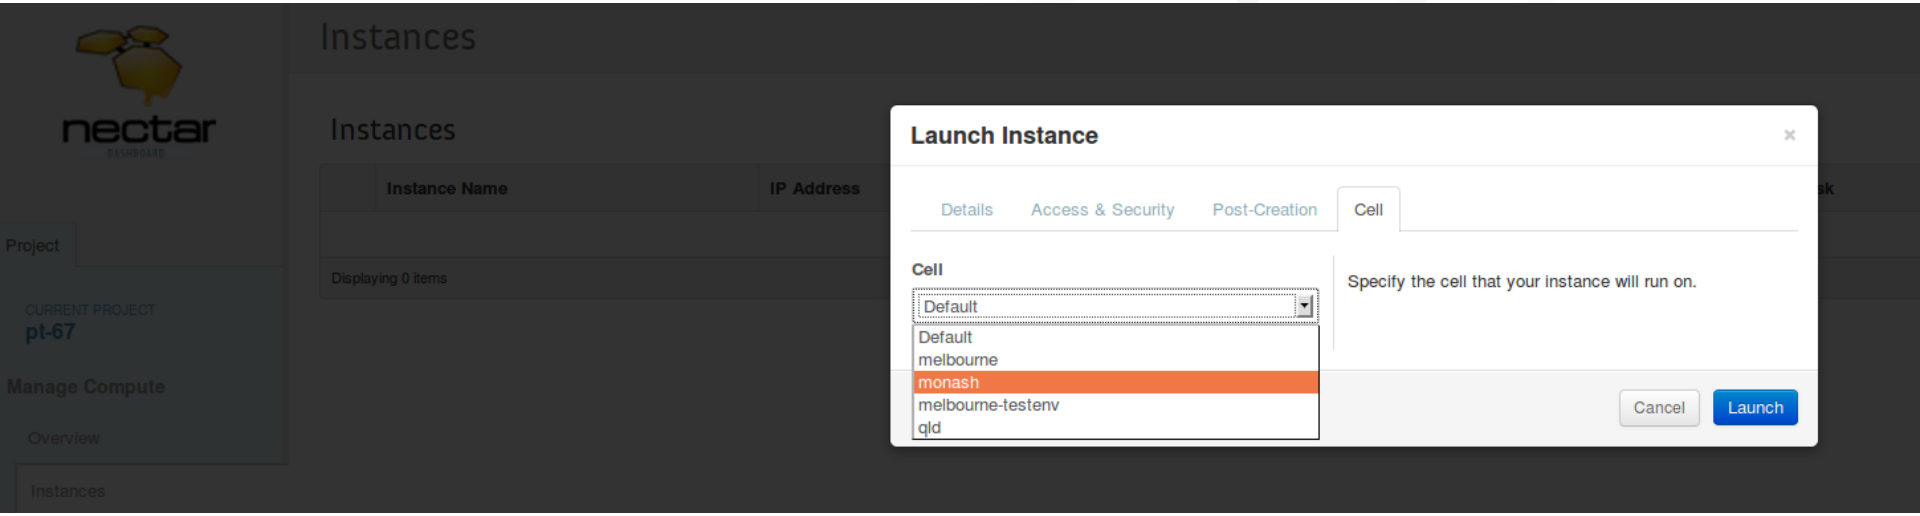

nectar

Cells - how we use it

- Single central API endpoint
- Each site has one or more cell
- Some sites have multiple datacenters = multiple cells
- Give sites more control
  - Different internal scheduling
  - Different compute node types, different flavours

nectar

Cells - early cells code

- Built on top of stable/folsom
- Uses an early version of cells code from Chris Behrens
  - Base of code is same as Grizzly cells
- Modified to support our use cases
  - Security group syncing
  - Allow users to select what cell they want
  - EC2 mapping syncing
  - Cell Scheduling
  - Continually rebasing from stable/folsom
  - Fixing bugs

- Goal: to get our code and master in sync
  - Stop having our custom code

nectar

Thanks to Chris Behrens

Cells and High Availability

- Need our central services to be reliable
- 98% uptime so not extreme
- Most pain points aren't with openstack itself
- Rabbit HA, Database replication
- Load balancing the APIs

Packaging and Patching

- Do a lot of this in house
- Can be slow to wait for it to go from master -> stable -> ubuntu package
- Package up own version of nova, horizon and keystone
- Follow stable branches in git with backports of certain things.

Packaging and Patching - Change Control

Try and copy openstack workflow for quality control

All | My | Projects | Groups | Documentation

Open Merged Abandoned

Tom Fifield <tom@tomfifield.net> | Settings | Sign Out

status:merged Search

Search for status:merged

	ID	Subject	Owner	Project	Branch	Updated
▸ ☆	I5e5ae594	add dell/mdraid checks, clean up old files (Merged)	Puppet Server	internal/puppet-physical	master (master)	Apr 11
☆	I137e1b30	move facts to physical module (Merged)	Puppet Server	internal/puppet-facts	master (master)	Apr 11
☆	If5933570	add mdadm fact (Merged)	Puppet Server	internal/puppet-facts	master (master)	Apr 11
☆	I0c0acac9	whitespace and formatting fixes (Merged)	Puppet Server	internal/puppet-iptables	master (master)	Apr 11
☆	I3c1f3586	clean up nagios class try 2 (Merged)	Puppet Server	NeCTAR-RC/puppet-nagios	master (master)	Apr 11
☆	Id40c2fd4	Add cells.scheduler_direct_only_cells option (Merged)	Kieran Spear	NeCTAR-RC/puppet-nova	production (production)	Apr 9
☆	Ief05150e	Use preproduction cell/role settings from hiera (Merged)	Kieran Spear	internal/puppet-dashboard	production (production)	Apr 9
☆	Icb172db0	fix ipmi user logic to cover different ipmi cards (Merged)	Puppet Server	internal/puppet-physical	master (master)	Apr 9
☆	Ief05150e	Use preproduction cell/role settings from hiera (Merged)	Puppet Server	internal/puppet-dashboard	master (preprod)	Apr 8
☆	I8e32dbad	updated ganglia module with extra stats (Merged)	Puppet Server	internal/puppet-ganglia	master (master)	Apr 8
☆	Idcefd385	use filename for conf file (Merged)	Puppet Server	internal/puppet-ganglia	master (master)	Apr 5
☆	I56f39869	simplify ganglia class (Merged)	Puppet Server	internal/puppet-ganglia	master (master)	Apr 5
☆	Ice91f9c2	nfs check using python module (Merged)	Puppet Server	internal/puppet-ganglia	master (master)	Apr 4
☆	Id40c2fd4	Add cells.scheduler_direct_only_cells option (Merged)	Kieran Spear	NeCTAR-RC/puppet-nova	master (add-direct-only-option)	Apr 4
☆	Ie113fa2d	Remove index view on support (Merged)	Sam Morrison	internal/puppet-apache	production (production)	Apr 4
☆	Ie113fa2d	Remove index view on support (Merged)	Puppet Server	internal/puppet-apache	master (master)	Apr 4
☆	I8d18424f	Added block migration flags to the nova.conf templates. (Merged)	Sam Morrison	NeCTAR-RC/puppet-nova	production (production)	Apr 3
☆	Ib3dfd076	Disable fixed ips quota (Merged)	Sam Morrison	NeCTAR-RC/puppet-nova	production (production)	Apr 3
☆	I8d18424f	Added block migration flags to the nova.conf templates. (Merged)	Jerico Revote	NeCTAR-RC/puppet-nova	master (master)	Apr 3
☆	Ib59af58b	Push out nagios check (Merged)	Sam Morrison	NeCTAR-RC/puppet-glance	production (production)	Mar 28
☆	I651f7bc5	typo (Merged)	Sam Morrison	NeCTAR-RC/puppet-swift	production (production)	Mar 28
☆	I27180a85	Push out nagios checks (Merged)	Sam Morrison	NeCTAR-RC/puppet-swift	production (production)	Mar 28
☆	I651f7bc5	typo (Merged)	Sam Morrison	NeCTAR-RC/puppet-swift	master (master)	Mar 28
☆	I27180a85	Push out nagios checks (Merged)	Sam Morrison	NeCTAR-RC/puppet-swift	master (master)	Mar 28

Packaging and Patching - Upgrades

- Started in production at Diablo
- Diablo -> Essex -> Folsom (current) -> Grizzly (in June)

- Zero downtime upgrades for instances
- Minimal downtime for central API services

- A lot of work
  - Database hacking
  - Lots of dry runs
  - Live migration needs to work

nectar

Looking to Grizzly

- Excited about increased support for operators
- Cells
- Identifying "good" images

- No quantum multi-host
- keystone db migrations

nectar

The Future - Ceilometer

- A better idea of how the cloud is being used
  - How much network are our scientists using?
  - Looking at whether scheduling and resource allocation is working
- Still seems hard to set up - but improving
- Important for sustainability - when the gov $ runs out

nectar

The Future - Other sites

- Currently 2 sites in production
- Queensland (north-eastern state) coming online in May
- 5 more to come

nectar

The Future - volumes

- Looking to ramp up volume offering significantly
- Might end up with many different vendor drivers
- Helps for people needing EC2 compatibility
- Will leverage funds from related project

nectar

National eResearch Collaboration Tools and Resources

Questions?

sam.morrison@unimelb.edu.au fifieldt@nectar.org.au

NeCTAR is an Australian Government project conducted as part of the Super Science initiative and financed by the Education Investment Fund. The University of Melbourne has been appointed the lead agent by the Commonwealth of Australia, Department of Innovation, Industry, Science and Research.

Objectives: to enhance research collaboration through the development of eResearch infrastructure.